Searching and Routing in Discrete and Continuous Domains

Prosenjit Bose (Carleton University),
Jean-Lou de Carufel (University of Ottawa),
Stephane Durocher (University of Manitoba),
Stefan Langerman (Université Libre de Bruxelles),
Ian Munro (University of Waterloo)

October 11, 2015–October 16, 2015

1 Overview of the Field

This workshop brought together researchers from the fields of Computational Geometry, Graph Theory, and other areas of mathematical research in Computer Science, working on various aspects of searching problems and routing problems in geometric settings. Workshop participants included a high representation of women researchers (13 of 32 attendees were women, corresponding to 40%) and researchers from Canada, the United States, Mexico, France, Belgium, Israel, Japan, Spain, Czech Republic, and the Netherlands.

In computer science, searching is the process of selecting a path in a given environment (the search space) to find a hidden target. The search space can be discrete (e.g., a graph) or continuous (e.g., a simple polygon and its interior). In computer science, routing is the process of selecting a path in a network along which to send network traffic (packets) from a given point of origin to a given destination point.

These two subjects are intimately related. In both, an agent (a searcher or a packet) is moving in some environment, trying to find a target (a hidden object or a destination). The agent must accomplish its task using only partial information. One of the main questions in these areas is actually to describe how much information the agent needs to be able to reach the target. In searching as well as in routing, a typical goal is to minimize the cost of the strategy, where the cost function can have different interpretations depending on the context. One cost function that has been widely considered in both settings is the competitive cost: the ratio between the cost of the strategy in the worst case and the cost of an optimal strategy obtained using complete information. Another similarity between these two fields of research is that the algorithms strongly depend on the geometry of the problems. Techniques used to solve problems in one topic can be applied to solve problems in the other topic.

In the last ten years, significant advances have been made on search and routing problems. This momentum was propelled by a number of factors:

1. Fundamental open problems in these areas were solved. For searching problems: the optimal strategy for searching on a line with or without extra information [9, 19]; the optimal strategy for searching on m concurrent rays with or without extra information [27]; the optimality of the spiral search in the plane [25]; searching for a line in the plane [2]. For routing problems: the classification of which Θ-graphs are spanners [1, 5, 10, 15, 16]; the optimal spanning ratio and optimal competitive cost of routing for infinite families of Θ-graphs [10]; several types of Delaunay triangulations have been proven to be spanners [8, 11]; moreover, local algorithms for these triangulations were designed [13, 14].

2. Search and routing problems captivated computer scientists for several reasons. No specialized mathematical or computer science background is required to understand the statements of the problems. Several fundamental questions in these areas can be solved in the most general frameworks. The algorithms we obtain are simple, though their proofs of correctness and optimality are intricate.

3. Another factor is the wealth of applications including motion planning, applications in biology including the study of predator and prey strategies [17] and partnership formation [24], network design, economy and game theory, operations research and management science.

The community of researchers in these areas has both matured and diversified in the last ten years. Four monographs [3, 4, 28, 29] and several survey papers have been released [16, 21, 22, 23, 30, 31]. A new generation of young scholars have done or are doing their graduate work in this area, in research groups in various academic centers. This workshop will be a great opportunity to bring together researchers working on various aspects of searching and routing problems. We believe that the discussions at the workshop will be fruitful and lead to new ideas and joint work in the future.

2 Highlights and Open Problems Discussed at the Workshop

The workshop was centred around seven talks, each focusing on open problems in some aspect of searching and routing. This set up sessions for participants to work in smaller groups on these open problems and related issues. This section presents an overview of the talks and the open problems presented in each.

2.1 Self-Approaching Graphs and Increasing Chord Graphs

Anna Lubiw, University of Waterloo, gave a presentation on this topic and introduced the following open problems.

A curve defined by a continuous function $f : [0, 1] \to \mathbb{R}^2$ is *self-approaching* if for all $0 \le a \le b \le c \le 1$ $\mathrm{dist}(b, c) \le \mathrm{dist}(a, c)$. The function f is also an *increasing chord* curve if if it is self-approaching in both directions, i.e., if for all $0 \le a \le b \le c \le d \le 1$ $\mathrm{dist}(b, c) \le \mathrm{dist}(a, d)$.

Local greedy routing is an algorithm that, given any pair of vertices s and t in a given straightline planar drawing of a graph G, identifies a path from s to t in G by iteratively constructing a route, selecting the next vertex adjacent to the endpoint that is closest to t, where distance usually refers to Euclidean distance.

1. Recognizing self-approaching/increasing chord drawings. This is a decision problem for which the input is a planar drawing of a graph.

   Open Problem 1 *Given a planar drawing of a graph G and vertices s and t in G, determine whether there is a self-approaching/increasing chord path from s to t in G.*

   If we can do this for any pair in polynomial time, then we can test every pair in polynomial time to determine whether G is self-approaching/increasing chord.

   Open Problem 2 *Characterize properties that must be true of a self-approaching/increasing chord graph G.*

2. Which graphs have increasing chord drawings? Which graphs have increasing chord drawings that support local greedy routing? This is a graph drawing problem for which the input is the combinatorial description of a given graph.

   Open Problem 3 *Do all 3-connected planar graphs have increasing-chord drawings?*

   Open Problem 4 *Do all 3-connected planar graphs have increasing-chord drawings that support local greedy routing?*

3. Given points in the plane, connect them with an increasing chord network. This is a graph drawing problem for which the input is a set of points in the plane.

   Open Problem 5 *Can every set of n points in the plane be connected by a straightline planar drawing that forms an increasing chord graph (without Steiner points)?*

   Open Problem 6 *Can every set of n points in the plane be connected by a straightline planar drawing that forms an increasing chord graph (without Steiner points) that supports local greedy routing?*

2.2 Compact Routing in Discrete Domains

Cyril Gavoille, University of Bordeaux, gave a presentation on this topic and introduced the following open problems.

Routing can also be achieved storing a routing table that stores precomputed routing information to assist with routing decisions in a given graph. The routing table can be distributed locally at vertices in the graph. An important cost is to measure the size of the largest routing table stored at any vertex. With limited table size a shortest path cannot be guaranteed. In this case the *stretch* (stretch factor) is measured, i.e., the worst-case ratio of the actual route length relative to the length of the shortest path. The goal is to find an efficient routing strategy that achieves a good trade-off between space and stretch on a given family of graphs. Both labelled and name-independent versions of the problem are considered. In the labelled problem, vertex labels can be used to encode partial information to assist with routing.

Open Problem 7 *With vertex labelling, is stretch $4k - O(1)$ achievable in $\tilde{O}(n^{1/k})$ space?*

The lower bound on stretch is $2k - 1$ for $k \in \{1, 2, 3, 5\}$.

Open Problem 8 *Do the bounds for the labelled and name-independent versions of the problem differ?*

When $k \in \{1, 2\}$ the bounds are equal.

Open Problem 9 *Do the bounds for the directed and undirected versions of the problem differ?*

Open Problem 10 *What is the best stretch achievable with polylogarithmic space in the name-independent setting on unweighted trees?*

The current best known stretch factor is 17.

Open Problem 11 *With vertex labelling on graphs with treewidth k can we achieve shortest path using $o(k \log^2 n)$-bit tables and labels?*

This property holds for trees and weighted outerplanar (and even $K_{2,4}$-minor free graphs): $\Theta(\log^2 n/ \log\log n)$ bits are necessary and sufficient. What about weighted series-parallel graphs?

Open Problem 12 *Give tight asymptotic bounds on the space required for shortest paths in planar graphs using polylogarithmic vertex labelling.*

Current bounds are $\Omega(n^{1/3})$ and $O(n)$, with the current best upper bound at $7.18n$ bits.

Open Problem 13 *Is $O(1)$-strech possible for general graphs using tables of size $\tilde{O}(\deg(u))$? What about in bounded-degree graphs?*

No lower bounds are known.

Open Problem 14 *Bound average stretch, ϵ-slack routing.*

The labelled and name-independent settings differ here.

Open Problem 15 *Give a distributed algorithm for constructing routing tables.*

2.3 Searching and Patrolling on a Line

Jean-Lou De Carufel, Ottawa University, gave a presentation on this topic and introduced the following open problems.

A line segment (fence) is to be patrolled by a given set of mobile guards, each of which moves with a given maximum speed, $v_1, \ldots, v_k$, such that each point on the fence is visited by some guard at least once per time unit. The *partition=based strategy* is to assign to each guard continuous interval of the fence to patrol by repeatedly visiting its left and right endpoints. Using this strategy, a fence of length $\frac{1}{2}\sum_{i=1}^{k} v_i$ can be patrolled.

Open Problem 16 *Given a set of k guards with maximum speeds $v_1, \ldots, v_k$, what is the longest fence that can be patrolled?*

Open Problem 17 *For which values of k is the partition-based strategy optimal?*

The strategy is known to be optimal for $k \leq 3$ and there exist instances of $k \geq 6$ for which it is known not to be optimal. I.e., for any $k \geq 6$ there exist $v_1, \ldots, v_k$ such that a fence of length greater than $\frac{1}{2}\sum_{i=1}^{k} v_i$ can be patrolled by k guards with respective maximum speeds $v_1, \ldots, v_k$. When $v_1 = \cdots = v_k$ then the partition-based strategy is known to be optimal for all k.

2.4 Searching for the Sink in Grids with Unique Sink Orientations

Luis Barba, Universite Libre de Bruxelles, gave a presentation on this topic and introduced the following open problems.

An oriented grid is a *unique sink orientation (USO)* if every subgrid has a unique sink. A sufficient condition is that this property holds for every 2×2 subgrid. If no 2×2 subgrid has two sinks and no 2×2 subgrid is a cycle, then the entire grid is acyclic. An *edge query* returns the orientation of the edge. A *vertex query* returns the set of edges adjancent to the vertex. How many queries are necessary to identify the sink vertex in a USO grid

Open Problem 18 *Give tight bounds on the number of edge/vertex queries necessary and sufficient to identify the sink vertex in a USO grid, in both the deterministic and randomized settings.*

Current bounds:

	deterministic	randomized
edge queries	$\Omega(n)$, $O(n^{1.58})$	$\Theta(n)$
vertex queries	$\Omega(n)$, $O(n \log n)$	$\Omega(\log n)$, $O(\log^2 n)$

Open Problem 19 *How many edge/vertex queries are necessary and sufficient to identify the sink vertex in a three-dimensional USO grid?*

2.5 Local Routing in Geometric Graphs

Stephane Durocher, University of Manitoba, gave a presentation on this topic and introduced the following open problems.

A *local routing algorithm* determines a sequence of forwarding decisions that define a route from a vertex s to a vertex t, where each internal node along the path selects one of its neighbours to extend the path as a function of its local network neighbourhood and limited information about the target node t. Stateless predecessor-oblivious local routing algorithms are known for triangulations. Every stateless predecessor-oblivious local routing algorithm is defeated by some convex subdivision. $O(\log n)$-bit predecessor-aware local routing algorithms are known for planar and near-planar graphs, and in more general classes of non-geometric graphs. The goal is to bridge the gap between local routing algorithms that use a single state bit and those that use $\Theta(\log n)$ state bits.

Open Problem 20 *On what classes of geometric graphs can local routing succeed using $o(\log n)$ state bits?*

Open Problem 21 *Can a predecessor-aware geometric local routing algorithm using $O(\log r)$ state bits succeed on planar subdivisions in which each face has at most r reflex vertices?*

Open Problem 22 *Can a geometric local routing algorithm using $O(1)$ state bits guarantee $o(n)$ competitive ratio (worst-case ratio of actual route length to shortest path length) on some class of geometric graphs (e.g., convex subdivisions)?*

Open Problem 23 *Can a stateless predecessor-oblivious local routing algorithm guarantee $O(1)$ competitive ratio on Θ-monotone graphs?*

Various greedy strategies guarantee delivery on Θ-monotone graphs because forwarding the message to any neighbour in the Gabriel circle from u to t (which must be non-empty) guarantees forward progress.

Open Problem 24 *Can a predecessor-oblivious local routing algorithm using $o(\log n)$ state bits succeed when all stateless predecessor-aware local routing algorithms fail (on some interesting class of geometric graphs)? I.e., is predecessor awareness always more powerful than $o(\log n)$ state bits?*

Open Problem 25 *Are $\Omega(\log n)$ state bits necessary for predecessor-aware local routing on planar subdivisions?*

2.6 Beacon-Based routing and Art Gallery Problems

Irina Kostitsyna, TU Eindhoven, gave a presentation on this topic and introduced the following open problems.

The attraction region of a point p in a polygon P is the locus of points q in P, such that there is an attraction path from q to p in P, where an attraction path travels directly toward p in the interior of P and travels along the boundary of P toward p while on the boundary of P. The inverse attraction region of a point q in P is the locus of points q in P such that there is an attraction path from q to p in P.

Open Problem 26 *What is the combinatorial complexity of attraction/inverse attraction regions for a given point p in a polygon q? How quickly can this region be computed?*

Current bounds for attraction region:

	simple polygon	polygon with h holes
complexity of region	$O(n)$	$O(n)$
computation time	$O(n \log n)$	$O(nh)$

Current bounds for inverse attraction region:

	simple polygon	polygon with holes
number of components	$\Omega(n)$	$\Omega(n^2)$
complexity of region	$\Omega(n)$	$\Omega(n^2)$
computation time	$O(n^3)$	$O(n^3)$

Beacon problems can be cast as art gallery problems: what is the minimum number of attraction points necessary to cover a given polygonal region? Both the attraction and inverse attraction are NP-hard.

Finally, beacons can be used for routing, where a beacon attracts a mobile agent m until m reaches the beacon, at which point the next beacon is activated, and so on until m reaches the target point. The objective is to minimize the number of beacons. The problem is known to be NP-hard in the all-to-all, s-to-all, and all-to-t settings.

Open Problem 27 *Given points s and t in a polygon P, what is the complexity of selecting a set of beacons of minimum cardinality to guide a mobile agent from s to t by attraction?*

2.7 Batched Point Location in SINR Diagrams via Algebraic Tools

Matya Katz, Ben-Gurion University, gave a presentation on this topic and introduced the following open problems.

The *signal to interference plus noise (SINR)* model in wireless networks introduces interference and power to provide a more realistic model for wireless communication. Given a set of wireless nodes positioned at points $p_1, \ldots, p_n$ in $\mathbb{R}^d$, each of which has an associated non-negative transmission power level $r(p_i)$, a receiver located a the point v can successfully receives a radio signal transmitted by the node p_i if and only if

$$\frac{r(p_i)}{|v-p_i|^\alpha} \geq \beta \sum_{j \neq i} \frac{r(p_j)}{|v-p_j|^\alpha} + N,$$

where $\alpha \geq 1$, $\beta \geq 1$, and $N \geq 0$ are fixed. Simpler versions of the model assign uniform power to the nodes, i.e., $r(p_i) = r(p_j)$ for all i and j, and set $N = 0$. The SINR model partitions $\mathbb{R}^d$ into cells, each of which is the locus of points that can successfully receive a transmission from p_i for some i. Between the cells are dead zones, in which the interference is too high to receive a transmission from any p_i.

Open Problem 28 *Find efficient algorithms for computing SINR the regions for a given set of wireless nodes and power levels.*

3 Scientific Progress and Outcome of the Meeting

The week was dominated by our breaking into working groups (typically at least two sessions per day) focusing on the open problems proposed in the talks and related problems. These groups were flexible and their composition changed on each meeting, depending on the expertise and interests of the participants. Interaction among participants has continued and multiple papers that resulted from collaborations at the workshop are currently in various stages of submission and review.

The open problems described in Section 2.1 led to two papers on "Strongly monotone drawings of planar Hamiltonian graphs" and "Gabriel Triangulations and Angle-Monotone Graphs: Local Routing and Recognition". These are in preparation for submission to the upcoming Symposium on Graph Drawing. The following results are included in the paper.

A path $v_0, v_1, \ldots, v_n$ in a graph G is angle-monotone with width γ and orientation β if every vector (v_i, v_{i+1}) lies in a cone of angle γ with orientation β. Notice that this implies that the whole path $v_i, \ldots, v_n$ lies in this cone.

1. In $O(mn + n^2 \log n)$ time we can determine if a graph G is angle-monotone with width γ, where m is the number of edges and n is the number of vertices of G.

2. Angle-monotone graphs with width γ are spanners with spanning ratio $1/\sin(\gamma/2)$. Given a set of points in the plane, we can always construct a plane angle-monotone graph with width $2\pi/3$. Prior to this work, it was not known if there existed plane angle-monotone graphs of bounded width that can be constructed on any point set.

3. Dehkordi, Frati, and Gudmundsson [18] showed that if the Delaunay triangulation of a point set has the property that no internal angle of a triangular face is more than $\pi/2$, then this graph is angle-monotone with width $\pi/2$. Note that this is a better width than the previous result, however, not all point sets admit such a Delaunay triangulation. We show how to locally route on these graphs (when they exists) such that the routing ratio is $1 + \sqrt{2}$ and prove that this bound is tight.

The open problems on patrolling described in Section 2.3 were also discussed extensively. Discussion focused on determining whethere there exists an optimal periodic strategy, and attempting to improve the best known results for a fixed number of agents (on the circle and on the line segment).

References

[1] O. Aichholzer, S.W. Bae, L. Barba, P. Bose, M. Korman, A. van Renssen, P. Taslakian, and S. Verdonschot, Θ-3 is connected, In *Proc. CCCG* (2013), 205–210.

[2] G. Aloupis, J. Iacono, J. Lenchner, and O. Ozkan, Locating a line at unit distance with multiple agents, In *P*roc. TJJCCGG(2012).

[3] S. Alpern, R. Fokkink, L. Gasieniec, R. Lindelauf, and V.S. Subrahmanian, Search theory: a game theoretic perspective, Springer (2013).

[4] S. Alpern and S. Gal, The theory of search games and rendezvous. International Series in Operations Research & Management Science, Springer (2003).

[5] R.A. Baeza-Yates, J.C. Culberson, and G.J.E. Rawlins, Searching in the plane, Inf. Comput., 106(2):234–252 (1993).

[6] L. Barba, P. Bose, J.-L. De Carufel, A. van Renssen, and S. Verdonschot, On the stretch factor of the Θ-4 graph, In *Proc. WADS* (2013) 109–120.

[7] L. Barba, P. Bose, M. Damian, R. Fagerberg, J. O'Rourke, A. van Renssen, P. Taslakian, and S. Verdonschot, New and improved spanning ratios for Yao graphs. ArXiv, 1307.5829 (2013).

[8] P. Bose, P. Carmi, S. Collette, and M. Smid, On the stretch factor of convex delaunay graphs. JoCG, 1(1):41–56 (2010).

[9] P. Bose, J.-L. De Carufel, and S. Durocher, Revisiting the problem of searching on a line, In *Proc. ESA*, 205–216 (2013).

[10] P. Bose, J.-L. De Carufel, P. Morin, A. van Renssen, and S. Verdonschot, Optimal bounds on Θ-graphs: more is not always better, In *Proc. CCCG*, 305–310 (2012).

[11] P. Bose and J.M. Keil, On the stretch factor of the constrained Delaunay triangulation. In ISVD, 25–31 (2006).

[12] P. Bose, A. Maheshwari, G. Narasimhan, M. Smid, and N. Zeh, Approximating geometric bottleneck shortest paths. Comput. Geom., 29(3):233–249 (2004).

[13] P. Bose and P. Morin. Competitive online routing in geometric graphs, Theor. Comput. Sci., 324(2-3):273–288 (2004).

[14] P. Bose and P. Morin. Online routing in triangulations, SIAM J. Comput., 33(4):937–951 (2004).

[15] P. Bose, P. Morin, A. van Renssen, and S. Verdonschot, The Θ-5 graph is a spanner. In *Proc. WG* (2013).

[16] P. Bose, A. van Renssen, and S. Verdonschot, On the spanning ratio of Θ-graphs. In *Proc. WADS*, 182–194 (2013).

[17] M. Broom, Interactions between searching predators and hidden prey, (S. Alpern, R. Fokkink, L. Gasieniec, R. Lindelauf, and V.S. Subrahmanian, eds.) Search theory: a game theoretic perspective, Springer (2013).

[18] H. R. Dehkordi, F. Frati, J. Gudmundsson, Increasing-chord graphs on point sets, J. Graph Algorithms Appl. 19(2):761–778 (2015).

[19] E. Demaine, S.P. Fekete, and S. Gal, Online searching with turn cost, Theor. Comput. Sci., 361(2-3):342–355 (2006).

[20] E. Demaine, J.B. Mitchell, and J. O'Rourke, Yao-Yao graph a spanner? (problem 70), In *The Open Problems Project* (2008).

[21] S. Gal, Search games, Wiley Encyclopedia of Operations Research and Management Science (L. A. Cox, Jr., ed) Wiley (2011).

[22] W. Huang, S. Chen, and W. Wang, Navigation in spatial networks: A survey, Physica A: Statistical Mechanics and its Applications (2013) in press.

[23] I. Kanj, Geometric spanners: recent results and open directions, In *Proc. ICCIT*, 78–82 (2013).

[24] S. Kinsella and D.M. Ramsey, A model of partnership formation with friction and multiple criteria, Search theory: a game theoretic perspective (S.Alpern, R. Fokkink, L. Gasieniec, Roy Lindelauf, and V.S. Subrahmanian, eds) Springer (2013).

[25] E. Langetepe, On the optimality of spiral search, In *Proc. SODA*, 1–12 (2010).

[26] M. Li, P.-J. Wan, Y. Wang, and O. Frieder, Sparse power efficient topology for wireless networks, In *Proc. HICSS*, 3839–3848 (2002).

[27] A. Lopez-Ortiz and S. Schuierer, The ultimate strategy to search on m rays? Theor. Comput. Sci., 261(2):267–295 (2001).

[28] G. Narasimhan and M. Smid, Geometric spanner networks, Cambridge University Press (2007).

[29] D. Peleg, Distributed computing: a locality-sensitive approach, Monographs on Discrete Mathematics and Applications, Society for Industrial and Applied Mathematics (2000).

[30] Y. Wang, Topology control for wireless sensor networks, Wireless Sensor Networks and Applications (Y. Li, M. Thai, and W. Wu, eds) 113–147, Springer (2008).

[31] W. Wu and Z. Zhang, Geometric optimization in wireless networks, in *Handbook of Combinatorial Optimization*, 1415–1457, Springer (2013).